INVESTMENTS
Lecture 5: More Dynamic Models

Philip H. Dybvig
Washington University in Saint Louis

- Traditional portfolio insurance
  - calculations
  - practical considerations
- constant proportions portfolio insurance

Portfolio insurance

Wouldn't it be nice if we could participate in the market when it goes up but be protected from losses when the market falls? No, this isn't wishful thinking or a something-for-nothing scheme. It is portfolio insurance, which trades some value on the upside for the guarantee on the downside. In terms of the market M_T, portfolio insurance offers a payoff

$$W_T = max\left(f, kW_0\frac{M_T}{M_0}\right)$$

where f is the floor (miniumum) payoff and k gives the amount of participation on the upside. Choosing $f = 0$ and $k = 1$ is simply investing in the market. Usually, we choose $f > 0$ and $1 > k > 0$: in this way we give up some participation on the upside in exchange for a floor on the downside.

Portfolio insurance: history

The portfolio insurance strategy was developed by Berkeley academics Mark Rubinstein and Hayne Leland, and was marketed by LOR (Leland-O'Brien-Rubinstein). Using a replicating strategy from option pricing theory, it is possible to create the desired terminal payoff through asset allocation between stocks and bonds without trading options.

Portfolio insurance: stock plus a put option

Recall that a *European put option* gives the owner the right (but not the obligation) to sell one share of the underlying security at a strike price of f. It is useful to interpret the portfolio insurance payoff as an underlying investment in equities plus a put option on that investment.

$$
\begin{aligned}
W_T &= max\left(f, kW_0\frac{M_T}{M_0}\right) \\
&= kW_0\frac{M_T}{M_0} + max\left(0, f - kW_0\frac{M_T}{M_0}\right).
\end{aligned}
$$

We can see that the portfolio insurance payoff is the same as investing kW_0 in the market and a put option on that investment with an exercise price of f.

Portfolio insurance: paying separately

In general, we need to make sure that our initial investment is correct, which means that the amount we pay for the put is just equal to the savings from choosing $k < 1$. This implies some sort of search for what values of k and f go together. A simpler approach is to think about paying for the portfolio insurance separately, in which case we take $k = 1$ and think of paying for the put out of a separate account.

To see how this works, consider an example with an initial investment of 100 and a floor of 50 in a 2-period binomial model with $u = 1.5$, $d = 0.5$, and $r = 1.0$. We assume that we pay for the insurance separately.

Paying separately: solution

At the terminal nodes, investing 100 in the stock itself would become 225 (uu), 75 (ud or du), or 25 (dd). The put has value only in the state dd, when it is worth 25. In this example, the state prices for up and down are both $1/2$. Therefore, we have the following valuation tree for the put:

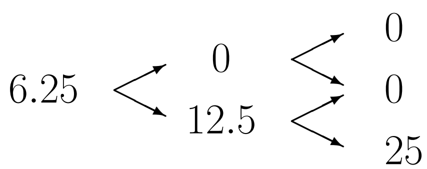

This corresponds to an initial payment of $25/4$ for the put and an initial investment of 100 in the underlying portfolio. The stock position is 100 in the underlying portfolio and $(0 - 12.5)/(1.5 - .5) = -12.5$ to replicate the put for an overall stock position of 87.5.

In-class exercise: paying separately

Set up a portfolio insurance strategy for a 2-period example with $u = 2$, $d = 1/2$, and $r = 1$. Take the floor equal to the initial investment, which is 100, and assume the insurance is paid for separately ($k = 1$). Then, compute the cost of the put and the optimal initial proportion to be invested in stocks.

Built-in fee

In the in-class exercise, we assumed that the portfolio insurance was to be purchased separately. Of course, it is usually the case that we know how much money to invest in total, not the amount to invest before paying for portfolio insurance. This means that we have to choose values of f and k that are compatible. The usual practice is to choose a value of the floor (often equal to the initial investment) and then perform a search over k to find the one that makes the investment in stock plus the put worth the amount you want to invest. For purposes of communication with the client, it may be appropriate to promise a bit less to allow for commissions, tracking error, etc.

In-class exercise: built-in fee

Set up a portfolio insurance strategy for a 2-period example with $u = 2$, $d = 1/2$, and $r = 1$. Take the initial investment (including the "fee") to be 100 and set the floor $f = 50$. Compute the proportion k of the initial investment to allocate to the underlying portfolio, the cost of the put, and the initial optimal proportion to be invested in stocks. (Hints: k is chosen to make the value of the overall portfolio insurance payoff equal to 100. The floor "guarantee" will only be relevant in the worst state when the stock goes down twice.)

Portfolio insurance: using the Black-Scholes model

Recall that the Black-Scholes model gives the call option price as

$$SN(x_1) - BN(x_2)$$

where S is the value of the underlying stock and B is the value of a discount bond with face equal the strike price and maturing with the option, $N(\cdot)$ is the cumulative normal distribution function, where

$$x_1 = \frac{log(S/B)}{\sqrt{s^2T}} + \frac{\sqrt{s^2T}}{2}$$

and

$$x_2 = \frac{log(S/B)}{\sqrt{s^2T}} - \frac{\sqrt{s^2T}}{2}$$

where T is the time-to-maturity of the option and s^2 is the variance per unit time of the underlying stock. If there is a constant interest rate r, then $B = Xexp(-rT)$ where X is the exercise (strike) price.

Black-Scholes replicating portfolio

It is easy to read the replicating portfolio for the option off the Black-Scholes formula, since the first term gives the amount to be invested in the stock, and the second part the amount to invest in the riskfree bond. Once the parameters are estimated (primarily the variance parameter), it is easy to compute the portfolio insurance strategy.

Portfolio insurance: nonconstant riskfree rate

In our examples, we have assumed a constant riskless rate. In practice, we want to have a riskless asset for the relevant maturity. If the maturity is one year (or less), then we can use the T-Bill maturing at a date near our horizon for the portfolio insurance as the riskless asset. For longer horizons, we can use Treasury STRIPs. In general, we can think of using the value of the riskless investment as the numeraire (unit of account) and do all the analysis in that way. Formally, then, the variance we use should be the variance per unit time of returns in S/B, but in practice this is not significantly different than the variance per unit time of returns in S.

Portfolio insurance: execution

During the crash, many funds with portfolio insurance did badly because they had trouble executing trades they needed to make to adjust their hedges as the market fell. For this reason, it is probably prudent for a portfolio insurer to maintain a position in futures options that will reduce the need for trading in general and will automatically adjust the risk exposure during a crash in particular. Buying out-of-the money puts on the market for this purpose is a popular strategy, although this may tend to make those puts a bit more expensive than we think they should be.

Using futures to reduce transaction costs

As the market moves and the stock position gets out of line with what the model prescribes, it is useful to use futures contracts to manage the required day-to-day changes in risk exposure. This is much less expensive than trading the underlying stocks all the time. Normally, we would want to replace the futures position with actual shares when it gets large, especially if we are long stocks and short futures which provide an imprecise hedge. Interestingly, we do not have to trade stocks at all (if our investment policy and guidelines permit), and it is possible to run a portfolio insurance program using just futures and bonds.

Sources of tracking error

While the replicating strategy from option pricing theory works remarkably well, it does not hold exactly. Transaction costs make it infeasible to maintain the optimal hedge continuously; in fact, it is optimal to trade rather infrequently since transaction costs are of first order and the cost of being improperly hedged is of second order. When the hedge is not exactly the theoretical one, then a price move in one direction means an increase or decrease in value compared to an ideal hedge. This is one source of tracking error.

Another source of tracking error is when volatility is different than what is expected to be on average. For portfolio insurance, high volatility, especially repeated large up and down moves in a net flat market, is a primary source of negative tracking error in portfolio insurance. In the industry, this is referred to as “whip-saw.”

Tracking error and continuation

Tracking error creates a minor problem in continuation, since literally speaking the original strategy is infeasible (since, given how the market has done, we have the wrong wealth). The appropriate solution is normally to keep the promise f fixed but vary k to be consistent with current wealth. This is the same as pretending the market had the performance that would correspond to current wealth.

Taxable accounts

Portfolio insurance is generally inappropriate for taxable accounts. The reason is that the damaging realization of gains that could be deferred is more important than any potential benefit of customizing returns.

Constant proportions portfolio insurance: background

If the variance of the market (or another risky portfolio) s^2 and the riskfree rate are constant, then we have that a portfolio strategy of investing at all times a fraction k in the market and a fraction $1-k$ in the riskfree asset will have a final payoff of approximately

$$W_T \approx W_0 \left(\frac{M_T}{M_0}\right)^k exp\left(\left((1-k)r + \frac{(k-k^2)s^2}{2}\right)T\right),$$

where W_0 is initial wealth and M_T/M_0 is the final value of the market portfolio per dollar invested initially. This result is exact in the Black-Scholes framework, and is approximately true in the standard binomial model.

Portfolio insurance: choice of horizon

It is a common practice to restart portfolio insurance once a year, using the start-of-year wealth as the promise (f) to be used at the end of the year. Unfortunately, this policy leads to a "jerky" strategy with a sudden jump in the asset mix when we enter the new year. As we might expect, that sort of sudden change is not efficient and implies poor diversification over time, as shown in my paper on "how to throw away a million dollars in the stock market." One solution is to set the horizon for the portfolio insurance equal to your investment horizon. Another solution appropriate to problems with consumption withdrawal, is to use my "ratcheting" strategy that was awarded the Common Fund Prize.

Constant proportions portfolio insurance

Constant proportions portfolio insurance is a strategy in which part of the portfolio is invested in the riskless asset and the remainder is invested in fixed proportions in the risky asset and the riskless asset. This gives a payoff of $P + c(M_T/M_0)^k$, where the constant c can be computed from the above expression. The simple form of the portfolio rule says that at t we invest $k(W_t - Pexp(-r(T-t)))$ in the market and the remainder in the riskless asset.

Related strategies include the general linear strategy linear in wealth whose payoff has 3 terms: constant, linear, and quadratic, and strategies that are piecewise sums of powers and can be priced by Black-Scholes.

General payoffs

In general, option pricing will compute the initial wealth and dynamic trading strategy to generate any payoff as a function of the market (and other financial asset payoffs) we want at the end. The only real constraint is that we do not choose a payoff that costs more to replicate than out initial wealth. Having the payoff depend on commodity prices (e.g. energy prices for a university with especially high heating or air conditioning costs) would be an example of why we may want to condition our final payoff on financial prices even when we do not have a view about whether they will go up or down.

Active strategies

All the analysis we have done was based on the assumption that we are trying to tailor the payoff and that we do not have a particular view about asset returns beyond what might be inferred from simple historical averages. The same tools can be used when we do have a view about values. In this case, we may wish to tailor claims to fit our views and option pricing is a good tool to use for this purpose. For example, suppose that we think a firm will do very well if the market as a whole goes up but that its performance is hard to predict if the market goes down. In this claim, we may want to tilt into a payoff which is the excess return of the security over the market payoff when the market is up, but which is worthless when the market is down. Once we decide on the mathematical formula representing this strategy, pricing the payoff and computing its replicating trading rule can be performed as a straightforward application of option pricing methodology.